

11-15-00

A

**UTILITY PATENT APPLICATION TRANSMITTAL
(Small Entity)***(Only for new nonprovisional applications under 37 CFR 1.53(b))*Docket No.
DE919990052US1

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTS**Box Patent Application
Washington, D.C. 20231**

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

IMPROVING AVAILABILITY AND SCALABILITY IN CLUSTERED APPLICATION SERVERS

and invented by:

Frank Leymann and Dieter RollerIf a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 31 pages and including the following:
 - a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications *(if applicable)*
 - c. ☐ Statement Regarding Federally-sponsored Research/Development *(if applicable)*
 - d. ☐ Reference to Microfiche Appendix *(if applicable)*
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings *(if drawings filed)*
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

**UTILITY PATENT APPLICATION TRANSMITTAL
(Small Entity)**

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
DE919990052US1

Total Pages in this Submission

Application Elements (Continued)

3. ☒ Drawing(s) *(when necessary as prescribed by 35 USC 113)*
a. ☐ Formal b. ☒ Informal Number of Sheets 3
4. ☒ Oath or Declaration
a. ☐ Newly executed *(original or copy)* ☒ Unexecuted
b. ☐ Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional application only)*
c. ☒ With Power of Attorney ☐ Without Power of Attorney
d. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference *(usable if Box 4b is checked)*
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under
Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.
6. ☐ Computer Program in Microfiche
7. ☐ Genetic Sequence Submission *(if applicable, all must be included)*
a. ☐ Paper Copy
b. ☐ Computer Readable Copy
c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. ☐ Assignment Papers *(cover sheet & documents)*
9. ☐ 37 CFR 3.73(b) Statement *(when there is an assignee)*
10. ☐ English Translation Document *(if applicable)*
11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☒ Certificate of Mailing
☐ First Class ☒ Express Mail *(Specify Label No.):* EK830786260US

**UTILITY PATENT APPLICATION TRANSMITTAL
(Small Entity)**

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
DE919990052US1

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. ☒ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. ☐ Small Entity Statement(s) - Specify Number of Statements Submitted: _____

17. ☐ Additional Enclosures *(please identify below):*

Request That Application Not Be Published Pursuant To 35 U.S.C. 122(b)(2)

18. ☐ Pursuant to 35 U.S.C. 122(b)(2), Applicant hereby requests that this patent application not be published pursuant to 35 U.S.C. 122(b)(1). Applicant hereby certifies that the invention disclosed in this application has not and will not be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication of applications 18 months after filing of the application.

Warning

An applicant who makes a request not to publish, but who subsequently files in a foreign country or under a multilateral international agreement specified in 35 U.S.C. 122(b)(2)(B)(i), must notify the Director of such filing not later than 45 days after the date of the filing of such foreign or international application. A failure of the applicant to provide such notice within the prescribed period shall result in the application being regarded as abandoned, unless it is shown to the satisfaction of the Director that the delay in submitting the notice was unintentional.

UTILITY PATENT APPLICATION TRANSMITTAL
(Small Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
DE919990052US1

Total Pages in this Submission

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	48	- 20 =	28	x \$9.00	\$252.00
Indep. Claims	3	- 3 =	0	x \$40.00	\$0.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$355.00
OTHER FEE (specify purpose) _____					\$0.00
TOTAL FILING FEE					\$607.00

- ☐ A check in the amount of _____ to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. **09-0463** as described below. A duplicate copy of this sheet is enclosed.
- ☒ Charge the amount of **\$607.00** as filing fee.
 - ☒ Credit any overpayment.
 - ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
 - ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Dated: **November 14, 2000**

Signature

Lynn L. Augspurger, Reg. 24,227
IBM Corporation
Intellectual Property Law
2455 South Road, P386
Poughkeepsie, NY 12601
(845) 433-1174
Fax: (845) 432-9786

CC:

Docket No: DE919990052US1

IMPROVING AVAILABILITY AND
SCALABILITY IN CLUSTERED
APPLICATION SERVERS

APPLICATION FOR
UNITED STATES LETTERS PATENT

Express Mail Label No: EK830786260US

Date of Deposit: November 14, 2000

I hereby certify that this correspon-
dence is being deposited with the
United States Postal Service as
Express Mail Post Office to Addressee
Service under 37 CFR 1.10 on the date
indicated above and is addressed to
Box Patent Application, Commissioner
of Patents and Trademarks, Washington,
D. C. 20231.

Susan L. Nelson



I N T E R N A T I O N A L B U S I N E S S M A C H I N E S
C O R P O R A T I O N

0041694400

Improving Availability and Scalability in Clustered Application Servers

Field of the Invention

The present invention relates to a method and means of increasing the availability within a multitude of application servers providing services to a multitude of application clients. More particularly, the invention relates to a method and means of increasing the availability and scalability by workload balancing and distribution of workloads to improve resource consumption.

Background of the Invention

Enterprises depend on the availability of the systems supporting their day-to-day operations. A system is called **available** if it is up and running and is producing correct results. In a narrow sense, **availability** of a system is the fraction of time it is available. MTBF denotes the mean time before failure of such a system is denoted, i.e. the average time a system is available before a failure occurs (this is the reliability of the system). Ideally, the availability of a system is 1. Today, a system can claim high availability if its availability is about 99.999% (it is called fault tolerant if its availability is about 99.99%). J. Gray and A. Reuter, "Transaction processing: Concepts and Techniques", San Mateo, CA: Morgan Kaufmann 1993 give further details on these aspects.

Availability of a certain system or application has at least two aspects: in a first, narrow significance it relates to the question, whether a certain system is active at all providing its services; in a second, wider significance it relates to the question, whether this service is provided in a timely fashion offering a sufficient responsiveness.

As outlined in further details by D. Loshin, "High performance computing demystified", Academic Press, Inc., 1994 and K. Hwang, Advanced computer architecture: Parallelism, Scalability, Programmability, PMcGraw-Hill, Inc., 1993 and J. Gray and A. Reuter, "Transaction processing: Concepts and Techniques", San Mateo, CA: Morgan Kaufmann 1993 one fundamental mechanism to improve availability is based on "redundancy": The availability of hardware is

improved by building clusters of machines and the availability of software is improved by running the same software in multiple address spaces.

With the advent of distributed systems, techniques have been invented which use two or more address spaces on different machines running the same software to improve availability (often called **active replication**). Further details on these aspects may be found in S. Mullender, "Distributed Systems", ACM Press, 1993. In using two or more address spaces on the same machine running the same software which gets its request from a shared input queue the technique of warm backups is generalized by the **hot pool** technique.

C. R. Gehr et al., "Dynamic Server Switching for Maximum Server Availability and Load Balancing", U.S. Pat. No. 5,828,847 teaches a dynamic server switching system relating to the narrow significance of availability as defined above. The dynamic server switching system maintains a static and predefined list (a kind of profile) in each client which identifies the primary server for that client and the preferred communication method as well as a hierarchy of successively secondary servers and communication method pairs. In the event that the client does not have requests served by the designated primary server or the designated communication method, the system traverses the list to ascertain the identity of the first available alternate server-communication method pair. This system enables a client to redirect requests from an unresponsive server to a predefined alternate server. In this manner, the system provides a reactive server switching for service availability.

In spite of improvements of availability in the narrow sense defined above, this teaching suffers from several shortcomings. Gehr's teaching provides a reactive response only in case a primary server could not be reached at all. There are no proactive elements which prevent a client requests service from a non-responsive server. As the list of primary and alternate servers is statically predefined, there may be situations in which no server could be found at all or in which a server is found not before several non-responsive alternate servers have been tested. Moreover, Gehr's teaching does not allow for a dynamic workload balancing improving the availability in the wider sense, i.e. the responsiveness. According to Gehr, different clients might be controlled by different

lists of servers, which allow for a rudimentary and static workload balancing as different clients might send their requests to different servers. In a highly dynamic, worldwide operating network situation, where clients and servers permanently enter or leave the network and where the access pattern to the servers may change from one moment to the next, Gehr's teaching to improve the responsiveness is not adequate.

Another area of technology to be mentioned is the area of Transaction Processing monitors (TP monitors). TP monitors have been invented more than three decades ago to make effective use of expensive system resources (J. Gray and A. Reuter, "Transaction processing: Concepts and Techniques", San Mateo, CA: Morgan Kaufmann 1993): Ever increasing numbers of users had to be supported by a system, and it turned out that native operating system functionality did not suffice to allow this. A TP monitor as a layer on top of the operating system manages system resources at a much finer granularity, assigns them with care, and only if needed and only for the duration needed. As a result, for one and the same machine and operating system, a given application can support orders of magnitudes of more users when implemented in a TP monitor than when implemented based on native operating system features.

The very complex and sophisticated TP monitor technology is primarily limited to a certain server only and thus does not solve the availability problem of a distributed network of application servers.

With the advent of distributed systems supporting middleware, object request brokers (be it a CORBA implementation, or DCOM, or based on the Java beans model) favor **commodity cluster environments**, i.e. environments which are composed out of relatively cheap hardware; refer for instance to G.F. Pfister, In search of clusters - 2nd edition (Prentice Hall PTR, 1998). In such environments, the service providing software components are simply replicated on multiple machines to ensure scalability. But this requires a mechanism to assign service requests to the various service providers ensuring the effective exploitation of the cluster resources. As a consequence, the middleware implementing the distributed system has to deal with similar problems as traditional TP monitors did before (by the way, this is one of the reasons why such

systems are considered as "TP monitor like systems" today). In fact, the middleware is representing a single and central TP monitor.

These "TP monitor like systems" are much too complicated to develop as well as to administer. Moreover they themselves consume a significant amount of processing resources.

- 5 Despite of all of this progress, further improvements are urgently required supporting enterprises in increasing the availability of their applications and allowing for instance for electronic business on a 7 (days) * 24 (hour) basis; due to the ubiquity of worldwide computer networks at any point in time somebody might have interest in accessing a certain application server.

Summary of the Invention

10 An object of the invention is to increase the availability of a multitude of application servers providing services to a multitude of application clients.

It is a further object of the invention to increase the availability by providing a technology, which is highly responsive to dynamic changes of the availability of individual application servers within the network and which supports a less centralized decision process for workload balancing.

15 The invention relates to a technology of workload balancing for improved availability within a multitude of applications-servers and a multitude of application-clients interconnected with said application-servers by a communication network.

A proposed method comprises a first-step, wherein an application-client is caching availability data of a subset of currently active application-servers as potential target application-servers.

- 20 The method comprises a second-step, wherein for execution of an application-request the application-client selects an application-server from the subset based on a load-balancing decision of the application-client as target application-server. Finally the application-request is sent to the target application-server.

of the application-client as target application-server. Finally the application-request is sent to the target application-server.

The proposed technology increases the availability and scalability of a multitude of application servers providing services to a multitude of application clients. The current invention provides a proactive technology as it incorporates in the workload balancing decisions only application servers which are active, i.e. which actually are providing their services to application clients. This prevents that a client generates erroneous request routings requesting service from non-responsive servers. Through caching of availability data of active servers, a dynamic technique with ongoing processing is suggested being highly responsive to dynamic network situation where clients and servers permanently enter or leave the network and where the access pattern to the servers may change from one moment to the next. Thus, the invention can accommodate hot plug-in of server machines into application clusters, thus further increasing the scalability of the environment. Complicated administration efforts to associate application clients with application servers are complete avoided. By teaching that the application clients are executing the load-balancing decisions a less centralized decision process for workload balancing is supported. Significant processing burden is removed from the servers, where, according to the state of the art, the workload balancing decisions would be performed and which typically build the primary bottleneck for processing resources. Caching only a subset of active application servers in the application clients allows for an efficient usage of the limited resources on the client side: depending on the size of the subset efficient load balancing decisions can be performed by a client. It does not prerequisite a workload management facility within the environment, which eventually also eases the implementation of application servers.

These and other objects will be apparent to one skilled in the art from the following drawings and detailed description of the invention.

Brief Description of the Drawings

Fig. 1 is a diagram reflecting the concepts of an application server, a hot pool, an application cluster and an application client;

Fig. 2 is an overview of the basic concepts of the current invention; and

Fig. 3 visualizes the "watchdog monitoring" approach of the current invention according to which the watchdogs are monitoring not only the current activity status of there associated application servers but also are monitoring the activity status of other watchdogs.

Description of the Preferred Embodiment

The present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system - or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods.

Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form.

If the current specification is referring to an application it may be a computer program of any nature not limited to any specific type or implementation. The terms application client and application server have to be understood from a logical point of view only relating to some type of "instance". These terms do not distinguish necessarily different address space or even different computer systems.

The current invention is assuming a certain communication path between application client and application server; this does not mean that the invention is limited to a certain communication paradigm.

Also if the current specification is referring to a "database" the term is to be understood in a wide sense comprising not only actual databases (like relational, hierarchical databases etc.) but also simple files and the like.

Introduction and Problem

Enterprises depend on the availability of the systems supporting their day-to-day operations. A system is called available if it is up and running and is producing correct results. In a narrow sense, the availability of a system is the fraction of time it is available. In a second, wider sense availability relates to the question, whether an application service is provided in a timely fashion offering a sufficient responsiveness.

In specific the current invention is relating to environments called "application cluster" based on the following concepts which are also depicted in Figure 1:

An **application server** (110, 111 or 112) is an executable implementing a collection of related services - for instance including access to some shared remote database (100). A **hot pool** (110, 111, 112) is a collection of address spaces each of which runs the same application server and each of these application servers receive requests from an input queue (125), which is shared between the hot pool members. By a **server machine** (101, 102 or 103) we mean a certain physical machine which hosts a hot pool of application servers. An **application cluster** (120) is a collection of servers which fail independently and each server hosts a hot pool of application servers of the same kind.

Applications (130) request services from application servers via application clients. An **application client** (131) is an executable which runs on the same machine as the application and

which communicates with a server on behalf of the application. If the communication between the application client and a server is based on (asynchronous) reliable message exchange, the application server is said to be message based. In what follows we do assume message based communication between application clients and application servers; of course the invention is not limited to the message based communication paradigm as other paradigms may be used instead. Consequently, an application client requests the performance of a certain service by sending a corresponding message into the input queue of a hot pool of associated application servers on a particular machine.

Traditional workload management is often server-centric and very sophisticated, thus, hard to implement: It takes the particular environment into account, i.e. it is often geared towards the underlying operating system or even hardware environment. Thus, the corresponding components and even the basic concepts are not portable. These server-centric approaches do not solve the availability problem of a distributed network of application servers. These "TP monitor like systems" are much too complicated to develop as well as to administer. Moreover they themselves consume a significant amount of processing resources.

As the basic approach the current invention is starting from the observation that in application cluster environments, a much less sophisticated workload management suffice because the used machine resources (commodity clusters) are usually much cheaper and thus, a request distribution which is to a certain degree uneven (or sub optimal) can be tolerated.

Without workload management performed by the system itself, each client submits requests to a predefined server machine, i.e. an administrator must assign a target machine to each client: This is a cumbersome administrative task, typically resulting in a bad resource exploitation. A typical example of this category is the teaching of Gehr et al. as discussed above.

Furthermore, adding new machines to the application cluster environment requires additional administrative actions. As the plugged-in machines will receive requests only once associated with at least one client an administrator would have to make corresponding administrative decisions.

But such an approach is inadequate in a highly dynamic network situation where clients and servers permanently enter or leave the network and where the access pattern to the servers may change from one moment to the next.

A desired solution to these problems should therefore provide a proactive technology which automatically determines application servers entering (hot plug-in) or leaving the cluster, which is responsive to sudden changes of the access pattern of the multitude of application clients and which provide a less centralized decisions process for the balancing problem.

Basic Idea of the Solution

The fundamental principle of the current invention is that (elementary but also sophisticated) load balancing can be performed at the client site to ease the implementation of application server, i.e. application functions; in an ongoing process the server site may provide availability data on application servers and may thus influence load balancing decisions indirectly, but is not required to do so. To be more precise, it is the application client that performs load balancing, not the application server and not the exploited communication subsystem (e.g. messaging middleware, RPC middleware etc.). Nevertheless, the server site may assist the application client in workload management. Therefore in view of more advanced embodiments of the current invention the current teaching could be described as an application server assisted workload balancing of the application clients. Based on this principle a less centralized decisions process for the balancing problem has been introduced.

The further features of the current invention are described by referring to overview given in Fig. 2.

Associating Server Subsets To Clients

For the purpose of associating a certain application client with a certain subset of application servers the application client (201) caches a chosen subset of server machines as potential targets for its service requests (202). As indicated by the term cache, this subset of application servers may be volatile in nature and may undergo permanent changes. Different mechanisms can be used

to determine such a subset for each application client which of course may be different for different clients:

1. Via profiling each application client can be associated with a fixed subset (203) of the multitude of all application servers. From this set the application client can make its choice of targets. This subset can be fixed for all requests, or it can vary from request to request. Of course the profile (203) may undergo changes while an application client is active; these changes will then be reflected in the cached subset (202).
2. The application client can retrieve (204) from a "cluster database" (205) the necessary topology information, i.e. which servers build the cluster (209) of active application servers (206, 207, 208). This topology lookup can be done when the application client is started, or it can be done for each request (before or after the request is issued) or any combination of this. Doing it not just at startup time allows to consider even servers that have been hot plugged in recently, i.e. in this case also application servers which just became active or inactive will be reflected (an aspect which will be discussed further below).

That (sub)set generally stores availability data of currently active application servers. In the simplest case only the identity of the currently active application servers may be comprised. But in other embodiments of the invention for instance also the processing power of the individual server and/or their current workload or other availability related data may be reflected.

If the server subset is based on a profiling approach as indicated in (1) above, different mechanisms can be used to determine the corresponding subset, for example:

- i. The subsets associated with each client can be evenly distributed across all application clients in (the assumed) case that all application clients produce (on the average) the same workload.
- ii. The subsets are determined according to the expected workload generated by a client; for instance application clients with a large expected workload could be associated with application servers with great processing power and/or application

clients with a large expected workload could be associated with subsets with no or a modest intersection only. Of course many other distribution strategies are possible.

If the server subset is based on topology lookup, as indicated in (2) above, a component called
 5 "server monitor" or "server spray" (210) is introduced. The server spray actually accepts lookup requests and returns subsets of servers. Moreover the server spray is monitoring (212) the activity status of the application servers within the cluster. Thus, the server spray can be viewed as a server assisted workload balancing feature for application clients. The server spray may be located on a separate application server (211) or may be executing on one of the application servers (206
 10 - 208), which also provide application services to the clients. In performing the server subsetting, the server spray can

- i. randomly choose a subset of (available) servers from the cluster database (called simple topology lookup). This approach attempts to generate a workload distribution on the multitude of application servers, which in the mean is the same on all application servers.
 15
- ii. apply its own load balancing algorithm to figure out the appropriate subset returned to the requesting application client (called load balanced topology lookup). In this case, load balancing is assisted by the server site; nevertheless, this does not impact the application clients' load balancing mechanism as it provides a decision space within the clients' load balancing mechanism is operating. Existing algorithms or components might be reused to implement load balanced topology lookup on the client side.
 20

Load balanced topology lookup might be considering the number of clients already assigned to a particular server, thus evenly distributing subsets of servers across application clients; this might
 25 be appropriate in case all application clients produce (on the average) the same workload. It might be more sophisticated by requiring for instance that each application client must indicate the

expected amount of workload to be generated (for instance by simply specifying a "load class" with the lookup request). The returned subset might be impacted by this specification as it may be based on a load-balancing decision of the server spray taking into account the indication of the expected amount of workload, the application-servers processing power and/or the application-servers work load.

Client-Based Workload Balancing

The following outline on client-based workload balancing is independent from the mechanism used to determine the cached server subset, i.e. it can be based on profiling, simple topology lookup, or load balanced topology lookup. It is assumed that the elements of a server subset are "weighted" to reflect processing power and/or any other relevant characteristics relating to the availability of the individual application servers; note, that the case of non-weighted servers is covered by assigning the same weight to all servers.

Before issuing an application request, an application client selects (213) a certain one of the application servers comprised by the cached subset (202) as target application sever to send (214) the request to. The basic approach of the current invention is that this selection is base on a load-balancing decision of said application client. This basically realizes client-based workload management.

Various approaches can be used for this load balancing decision. For instance the application client could select the target application server randomly from the cached subset. If the availability data comprised by the subset stores the "weight" of the servers, for instance expressing the processing power and/or the application servers work load (for instance expressed by the depth of the input queue, last observed response time etc.), much more sophisticated load balancing methods can be used on the client side. Whatever load balancing method is exploited it is preferably utilized for approximating a even distribution of workload within said subset of application servers.

In doing so, each application client supports the even distribution of requests across all application servers available in the application cluster. Load balanced topology lookup might even amplify this effect. Ensuring an even distribution of requests results in the exploitation of all available resources in the overall environment, i.e. in scalability within the available boundaries.

5 Hot Plug-In Of Application Servers

Scalability is improved via hot plug-in of additional servers (i.e. application servers, which became active, or servers, which became inactive by have been removed from service) and delivering the information about new servers to all or selective application clients. The latter information can either be "pushed" (i.e. the server spray is taking initiative to send that information) to the clients or "pulled" (i.e. the application clients explicitly are requesting that information) by them.

"Pushing" this information in a single burst to all clients might swamp the network severely impacting the network throughput of "real" requests. This impacts the scalability of the overall environment and is thus avoided by the current invention. Instead, the push approach suggested here might be called **lazy multi-casting**, meaning that the server spray selects iteratively (with a time delay) subsets of application clients to which information about the newly added or removed server is sent. Again, much sophistication can be applied in determining the proper subsets of application clients and sequencing of the submission of this information.

Pulling this information requires that application clients perform from time to time a topology lookup (204): Such a lookup can be done whenever a client request has to be processed (before or after sending an application request), periodically, or randomly. This ensures that whenever a new server got added to the application cluster after a certain period of time it will receive service requests. In pulling topology information to refresh the server subset in its cache, an application client fundamentally supports scalability via hot plug-in of servers in commodity clusters.

When a new server machine is plugged in, the server spray could select collections of logged-in clients in a random fashion: It withdraws from their cached server subset randomly chosen servers and either pushes the corresponding update into their set or waits until they pull for topology

information. If the mechanism used does not store persistently for each application client its cached server subset, the application client will ignore the withdrawal if the server is not within the cached set of the application client. Again, to avoid network swamping, only a few application clients at a time get the modifications to their server subset pushed.

5 Maintaining The Set Of Available Servers

Workload management functions consider the amount of resources available on each server for processing client requests. On each server, this amount can be perceived as a continuous function in time of course depending on the status of work formerly submitted to the server. But these functions typically assume that the subject server is up and running in a high availability sense: if the server failed, these function may produce incorrect results based on the historical status and the inherent assumption of "always being up and running". I.e. the workload function might return a value from the interval $]0,1[$ although the real value of available resources is "0" because the server is not available at all, i.e. it does not provide any computing resources. Thus, by adding fault detection workload management will become more precise.

Each hot pool on a server of an application cluster is monitored by a "watchdog". Such a **watchdog** detects failed hot pool members and recreates them immediately in an automatic manner; this is depicted in Fig. 3. The current invention in addition suggests that the collection of watchdogs in an application cluster monitor themselves ("watchdog watchdogging" or watchdog monitoring). This is achieved via a communication medium that allows to run a protocol to receive information about the state of a hot pool on a given machine; for instance "heartbeat" protocols could be used for that purpose. When one of the server of the application cluster fails the watchdog on this server will be detected to have failed (with a latency that depends on the monitoring protocol used) by other watchdogs. As a result, the remaining watchdogs will be able to multi-cast this fact and send this information to the server spray (which will maintain the cluster database accordingly) so that the workload functions can make use of this information avoiding to schedule requests to the failed server. Similarly, when the server is available again the watchdogs will detect that. Once the workload functions got that information new requests can be submitted to the recreated server. The proposed approach significantly improves the responsiveness of the

cluster as it is avoided that application requests are sent to non-responsive servers; moreover processing resources are detected as soon as they become available and are included in the workload balancing process.

Thus according to the "watchdog monitoring" approach of the current invention, the watchdogs are monitoring not only the current activity status of there associated application servers but also are monitoring the activity status of other watchdogs.

Advantages of the Invention

The proposed technology increases the availability and scalability of a multitude of application servers providing services to a multitude of application clients. The current invention is providing a proactive technology as it incorporates in the workload balancing decisions only application servers which are active, i.e. which actually are providing their services to application clients. This prevents that a client generates erroneous request routings requesting service from non-responsive servers. Through caching of availability data of active servers a dynamic technique and ongoing process is suggested being highly responsive to dynamic network situation where clients and servers permanently enter or leave the network and where the access pattern to the servers may change from one moment to the next. Thus the invention can accommodate hot plug-in of server machines into application clusters further increasing the scalability of the environment.

Complicated or due its sheer complexity impossible administration efforts to associate application clients with application servers are complete avoided. By teaching that the application clients are executing the load-balancing decisions a less centralized decision process for workload balancing is supported. Significant processing burden is removed from the servers, where according to the state of the art the workload balancing decisions would be performed and which typically build the primary bottleneck for processing resources. Caching only a subset of active application servers in the application clients allows for an efficient usage of the limited resources (in terms of storage and processing power) on the client side: depending on the size of the subset efficient load balancing decisions can be performed by a client. Due to the teachings independence from the concrete algorithms for workload balancing the reuse of existing general purpose load balancing algorithms are possible. The proposed concepts are portable in the sense that its mechanisms do

not require hardware or software specific implementations (like highly specialized TPA Monitors do). As a result, it ensures a common behavior across environments. It does not prerequisite a workload management facility within the environment, which eventually also eases the implementation of application servers.

- 5 According to a certain embodiment of the invention, a single cluster database suffices as the only centralized feature, which in addition can be implemented easily with state of the art technology.

An embodiment wherein the active application servers are distributed approximately evenly over the various subsets of the various application clients turns out to be an efficient approach: such subsets can be determined performance-efficiently without increasing workload considerably.

- 10 Already an approximately even distribution of the application servers ensures that the probability that a servers is overrun is significantly reduced, that the overall throughput to the cluster is improved, that availability and responsiveness of the application servers is improved and that at the same time also improved scalability is offered.

- 15 The further embodiment of the invention introducing a server monitor or "server spray" simplifies the implementation of the application client. A mere retrieve request suffices to determine the subset, complicate database on the client side accesses are avoided. Finally the more resource expensive monitoring and administering of the complete set of application servers can be executed outside an application client on a more powerful system; thus the application clients have to cope with the much smaller subset only.

- 20 As a further advantage the server monitor could offer additional assistance if the subset of application servers provided is already the result of a load-balancing decision. This approach allows for a combination of global workload balancing, performed through the server monitor by prescribing to the application clients a "workload control space" (the subset), and a decentralized workload balancing, performed by the individual application clients balancing issued requests
- 25 within its assigned subset.

An embodiment of the invention wherein an application client sends an indication of the expected amount of workload to be generated to the server monitor permits the server monitor to make a more precise workload balanced subset selection leading to improved workload balancing decisions. Or on the other hand on the client side less sophisticated workload balancing methodologies can be applied without effecting the overall availability and scalability.

As further advantage the application, client may perform any type of workload balancing on the subset of application servers which may be adapted to the client's processing capabilities. This may range from a very economic but nevertheless efficient random selection to a more sophisticated workload balancing decision exploiting the application servers processing power, its current workload or its observed response time.

A further significant advantage is that the current teaching supports various degrees of the actuality of the availability data of the cached subset. It may range from a random or periodically actualization up to an actualization per request. Eventually the lazy-multi-cast approach combines high actuality of the availability data of the cached subset with a modest load on the network.

The introduction of watchdogs monitoring the activity status of its corresponding local application server which then inform the server monitor, guarantees a decentralized and thus efficient monitoring approach. By allowing in addition a watchdog to check the activity of other watchdogs, the problem of "checking the checker" is solved.

While the preferred embodiment of the invention has been illustrated and described herein, it is to be understood that the invention is not limited to the precise construction herein disclosed, and the right is reserved to all changes and modifications coming within the scope of the invention as defined in the appended claims.

Claims:

What is claimed is:

1. A computerized method of workload balancing for improved availability within a multitude of applications-servers and a multitude of application-clients interconnected with said application-servers by a communication network, said method comprising
caching within an application-client, availability data of a subset of currently active application-servers as potential target application-servers, and
selecting by execution of an application-request by said application-client, an application-server from said subset based on a load-balancing decision of said application-client as target application-server and is sending said application-request to said target application-server.
2. A method of workload balancing according to claim 1 further comprising,
retrieving by said application-client, said subset of currently active application-servers from a cluster database, said cluster database storing availability data of currently active application-servers.
3. A method of workload balancing according to claim 2 further comprising,
choosing said subset of currently active application-servers by approximating an even distribution of said application-servers with respect to other subsets of other application-clients.
4. A method of workload balancing according to claim 2 further comprising,

retrieving said subset of currently active application-servers by issuing a retrieve request to a server-monitor, said server-monitor monitoring said application-servers activity status.

5. A method of workload balancing according to claim 4 further comprising,

choosing by said server-monitor, said subset of currently active application-servers from said cluster database by a random selection.

6. A method of workload balancing according to claim 4 further comprising,

choosing by said server-monitor, said subset of currently active application-servers from said cluster database based on a load-balancing decision of said server-monitor.

7. A method of workload balancing according to claim 6 further comprising,

choosing by said server-monitor, said subset of currently active application-servers by approximating an even distribution of said application-servers with respect to other subsets of other application-clients.

8. A method of workload balancing according to claim 4 further comprising,

sending by said application-client together with said retrieve request, an indication of the expected amount of workload to be generated by said application-client, and

basing said server-monitor's load-balancing decision on said indication of the expected amount of workload and/or said application-servers processing power and/or said application-servers work load.

9. A method of workload balancing according to claim 1,

2 wherein selecting by said application-client comprises selecting said target
3 application-server randomly.

1 10. A method of workload balancing according to claim 1,

2 wherein said availability data of said subset of currently active application-servers
3 comprises said application-servers processing power and in addition said
4 application-servers work load and/or said application-servers observed response time and
5 wherein selecting by said application-client comprises selecting said target
6 application-server approximating an even distribution of workload within said subset of
7 currently active application-servers.

11. A method of workload balancing according to claim 1 further comprising,

2 dynamically updating said availability data of said subset of currently active
3 application-servers by explicitly requesting from said server-monitor an update of said
4 subset of currently active application-servers, and/or

5 dynamically updating by said server-monitor, said availability data of said subset of
6 currently active application-servers by receiving an update of said subset of currently active
7 application-servers.

1 12. A method of workload balancing according to claim 11,

2 wherein said explicitly requesting from said server-monitor an update of said subset of
3 currently active application-servers occurs at one of:

4 prior to an application-request, or
5 after sending an application-request, or
6 periodically, or
7 randomly.

1 13. A method of workload balancing according to claim 11,

2 wherein, in the event an additional application-server has become active or if an active
3 application-server has become inactive, said server-monitor takes the initiative of sending
4 said update of said subset of currently active application-servers.

1 14. A method of workload balancing according to claim 13,

2 wherein said server-monitor sending said update of said subset of currently active
3 application-servers comprises a lazy-multi-cast including;

4 iteratively selecting with a time delay, a collection of application-clients for said multitude
5 of application-clients, and

6 sending each application-client of said collection an update of said subse of currently active
7 application-servers.

1 15. A method of workload balancing according to claim 4,

2 wherein said application-servers each comprise a hot pool of one or a multitude of
3 application-instances, said application-servers being executed in the same or a multitude of
4 address-spaces, and

5 wherein for each of said hot pools a watchdog is monitoring said hot pool's activity status,
6 and

7 wherein each of said watchdogs is informing said server-monitor on the workload of said
8 monitored application-servers and its current activity status.

- 1 16. A method of workload balancing according to claim 15,
2 wherein at least one of said watchdogs is monitoring the activity status of other watchdogs
3 and, in the event that said one watchdog detects the unavailability of any one of said other
4 watchdogs, said one watchdog informs said server-monitor of said corresponding activity
5 status and said server-monitor updates said cluster database accordingly.
- 1 17. An apparatus for workload balancing giving improved availability within a multitude of
2 applications-servers and a multitude of application-clients interconnected with said
3 application-servers by a communication network, said apparatus comprising
4 means for caching within an application-client, availability data of a subset of currently
5 active application-servers as potential target application-servers, and
6 means for selecting by execution of an application-request by said application-client, an
7 application-server from said subset based on a load-balancing decision of said
8 application-client as target application-server and is sending said application-request to said
9 target application-server.
- 10 18. An apparatus for workload balancing according to claim 17 further comprising,
2 means for retrieving by said application-client, said subset of currently active
3 application-servers from a cluster database, said cluster database storing availability data of
4 currently active application-servers.
- 1 19. An apparatus for workload balancing according to claim 18 further comprising,
2 means for choosing said subset of currently active application-servers by approximating an
3 even distribution of said application-servers with respect to other subsets of other
4 application-clients.

- 1 20. An apparatus for workload balancing according to claim 18 further comprising,
2 means for retrieving said subset of currently active application-servers by issuing a retrieve
3 request to a server-monitor, said server-monitor monitoring said application-servers
4 activity status.
- 1 21. An apparatus for workload balancing according to claim 20 further comprising,
2 means for choosing by said server-monitor, said subset of currently active
3 application-servers from said cluster database by a random selection.
22. An apparatus for workload balancing according to claim 20,
2 means for choosing by said server-monitor, said subset of currently active
3 application-servers from said cluster database based on a load-balancing decision of said
4 server-monitor.
23. An apparatus for workload balancing according to claim 22 further comprising,
2 means for choosing by said server-monitor, said subset of currently active
3 application-servers by approximating an even distribution of said application-servers with
4 respect to other subsets of other application-clients.
- 1 24. An apparatus for workload balancing according to claim 20 further comprising,
2 means for sending by said application-client together with said retrieve request, an
3 indication of the expected amount of workload to be generated by said application-client,
4 and

means for basing said server-monitor's load-balancing decision on said indication of the expected amount of workload and/or said application-servers processing power and/or said application-servers work load.

25. An apparatus for workload balancing according to claim 17,

wherein said means for selecting by said application-client comprises means for selecting said target application-server randomly.

26. An apparatus for workload balancing according to claim 17,

wherein said availability data of said subset of currently active application-servers comprises said application-servers processing power and in addition said application-servers work load and/or said application-servers observed response time and wherein means for selecting by said application-client comprises means for selecting said target application-server approximating an even distribution of workload within said subset of currently active application-servers.

27. An apparatus for workload balancing according to claim 17 further comprising,

means for dynamically updating said availability data of said subset of currently active application-servers by explicitly requesting from said server-monitor an update of said subset of currently active application-servers, and/or

means for dynamically updating by said server-monitor, said availability data of said subset of currently active application-servers by receiving an update of said subset of currently active application-servers.

28. An apparatus for workload balancing according to claim 27,

wherein said explicitly requesting from said server-monitor an update of said subset of currently active application-servers occurs at one of:

prior to an application-request, or
after sending an application-request, or
periodically, or
randomly.

29. An apparatus for workload balancing according to claim 27,

wherein, in the event an additional application-server has become active or if an active application-server has become inactive, said server-monitor takes the initiative of sending said update of said subset of currently active application-servers.

30. An apparatus for workload balancing according to claim 29,

wherein said server-monitor sending said update of said subset of currently active application-servers comprises a lazy-multi-cast including;

means for iteratively selecting with a time delay, a collection of application-clients for said multitude of application-clients, and

means for sending each application-client of said collection an update of said subse of currently active application-servers.

31. An apparatus for workload balancing according to claim 20,

wherein said application-servers each comprise a hot pool of one or a multitude of application-instances, said application-servers being executed in the same or a multitude of address-spaces, and

wherein for each of said hot pools a watchdog is monitoring said hot pool's activity status,
and

wherein each of said watchdogs is informing said server-monitor on the workload of said
monitored application-servers and its current activity status.

32. An apparatus for workload balancing according to claim 31,

wherein at least one of said watchdogs is monitoring the activity status of other watchdogs
and, in the event that said one watchdog detects the unavailability of any one of said other
watchdogs, said one watchdog informs said server-monitor of said corresponding activity
status and said server-monitor updates said cluster database accordingly.

33. A computer program product comprising a computer usable medium having computer
readable program code means therein for workload balancing giving improved availability
within a multitude of applications-servers and a multitude of application-clients
interconnected with said application-servers by a communication network, said computer
program product comprising

computer readable program code means for caching within an application-client,
availability data of a subset of currently active application-servers as potential target
application-servers, and

computer readable program code means for selecting by execution of an
application-request by said application-client, an application-server from said subset based
on a load-balancing decision of said application-client as target application-server and is
sending said application-request to said target application-server.

34. The computer program product for workload balancing according to claim 33 further
comprising,

computer readable program code means for retrieving by said application-client, said subset of currently active application-servers from a cluster database, said cluster database storing availability data of currently active application-servers.

35. The computer program product for workload balancing according to claim 34 further comprising,

computer readable program code means for choosing said subset of currently active application-servers by approximating an even distribution of said application-servers with respect to other subsets of other application-clients.

36. The computer program product for workload balancing according to claim 34 further comprising,

computer readable program code means for retrieving said subset of currently active application-servers by issuing a retrieve request to a server-monitor, said server-monitor monitoring said application-servers activity status.

37. The computer program product for workload balancing according to claim 36 further comprising,

computer readable program code means for choosing by said server-monitor, said subset of currently active application-servers from said cluster database by a random selection.

38. The computer program product for workload balancing according to claim 36,

computer readable program code means for choosing by said server-monitor, said subset of currently active application-servers from said cluster database based on a load-balancing decision of said server-monitor.

1 39. The computer program product for workload balancing according to claim 38 further
2 comprising,

3 computer readable program code means for choosing by said server-monitor, said subset
4 of currently active application-servers by approximating an even distribution of said
5 application-servers with respect to other subsets of other application-clients.

1 40. The computer program product for workload balancing according to claim 36 further
2 comprising,

3 computer readable program code means for sending by said application-client together
4 with said retrieve request, an indication of the expected amount of workload to be
5 generated by said application-client, and

6 computer readable program code means for basing said server-monitor's load-balancing
7 decision on said indication of the expected amount of workload and/or said
8 application-servers processing power and/or said application-servers work load.

9 41. The computer program product for workload balancing according to claim 33,

2 wherein said computer readable program code means for selecting by said
3 application-client comprises computer readable program code means for selecting said
4 target application-server randomly.

1 42. The computer program product for workload balancing according to claim 33,

2 wherein said availability data of said subset of currently active application-servers
3 comprises said application-servers processing power and in addition said
4 application-servers work load and/or said application-servers observed response time and

wherein computer readable program code means for selecting by said application-client comprises computer readable program code means for selecting said target application-server approximating an even distribution of workload within said subset of currently active application-servers.

43. The computer program product for workload balancing according to claim 33 further comprising,

computer readable program code means for dynamically updating said availability data of said subset of currently active application-servers by explicitly requesting from said server-monitor an update of said subset of currently active application-servers, and/or

computer readable program code means for dynamically updating by said server-monitor, said availability data of said subset of currently active application-servers by receiving an update of said subset of currently active application-servers.

44. The computer program product for workload balancing according to claim 43,

wherein said explicitly requesting from said server-monitor an update of said subset of currently active application-servers occurs at one of:

prior to an application-request, or
after sending an application-request, or
periodically, or
randomly.

45. The computer program product for workload balancing according to claim 43,

wherein, in the event an additional application-server has become active or if an active application-server has become inactive, said server-monitor takes the initiative of sending said update of said subset of currently active application-servers.

- 1 46. The computer program product for workload balancing according to claim 45,
2 wherein said server-monitor sending said update of said subset of currently active
3 application-servers comprises a lazy-multi-cast including;
4 computer readable program code means for iteratively selecting with a time delay, a
5 collection of application-clients for said multitude of application-clients, and
6 computer readable program code means for sending each application-client of said
7 collection an update of said subse of currently active application-servers.
- 8 47. The computer program product for workload balancing according to claim 36,
9 wherein said application-servers each comprise a hot pool of one or a multitude of
10 application-instances, said application-servers being executed in the same or a multitude of
11 address-spaces, and
12 wherein for each of said hot pools a watchdog is monitoring said hot pool's activity status,
13 and
14 wherein each of said watchdogs is informing said server-monitor on the workload of said
15 monitored application-servers and its current activity status.
- 16 48. The computer program product for workload balancing according to claim 47,
17 wherein at least one of said watchdogs is monitoring the activity status of other watchdogs
18 and, in the event that said one watchdog detects the unavailability of any one of said other
19 watchdogs, said one watchdog informs said server-monitor of said corresponding activity
20 status and said server-monitor updates said cluster database accordingly.

**Improving Availability and Scalability
in Clustered Application Servers**

Abstract of the Disclosure:

The invention relates to a technology of workload balancing for improved availability within a
5 multitude of applications-servers and a multitude of application-clients interconnected with said
application-servers by a communication network. A proposed method comprises a first-step,
wherein an application-client is caching availability data of a subset of currently active
application-servers as potential target application-servers. The method comprises a second-step,
10 wherein for execution of an application-request the application-client selects an application-server
from the subset based on a load-balancing decision of the application-client as target
application-server. Finally the application-request is sent to the target application-server.




FIG. 1




FIG. 2

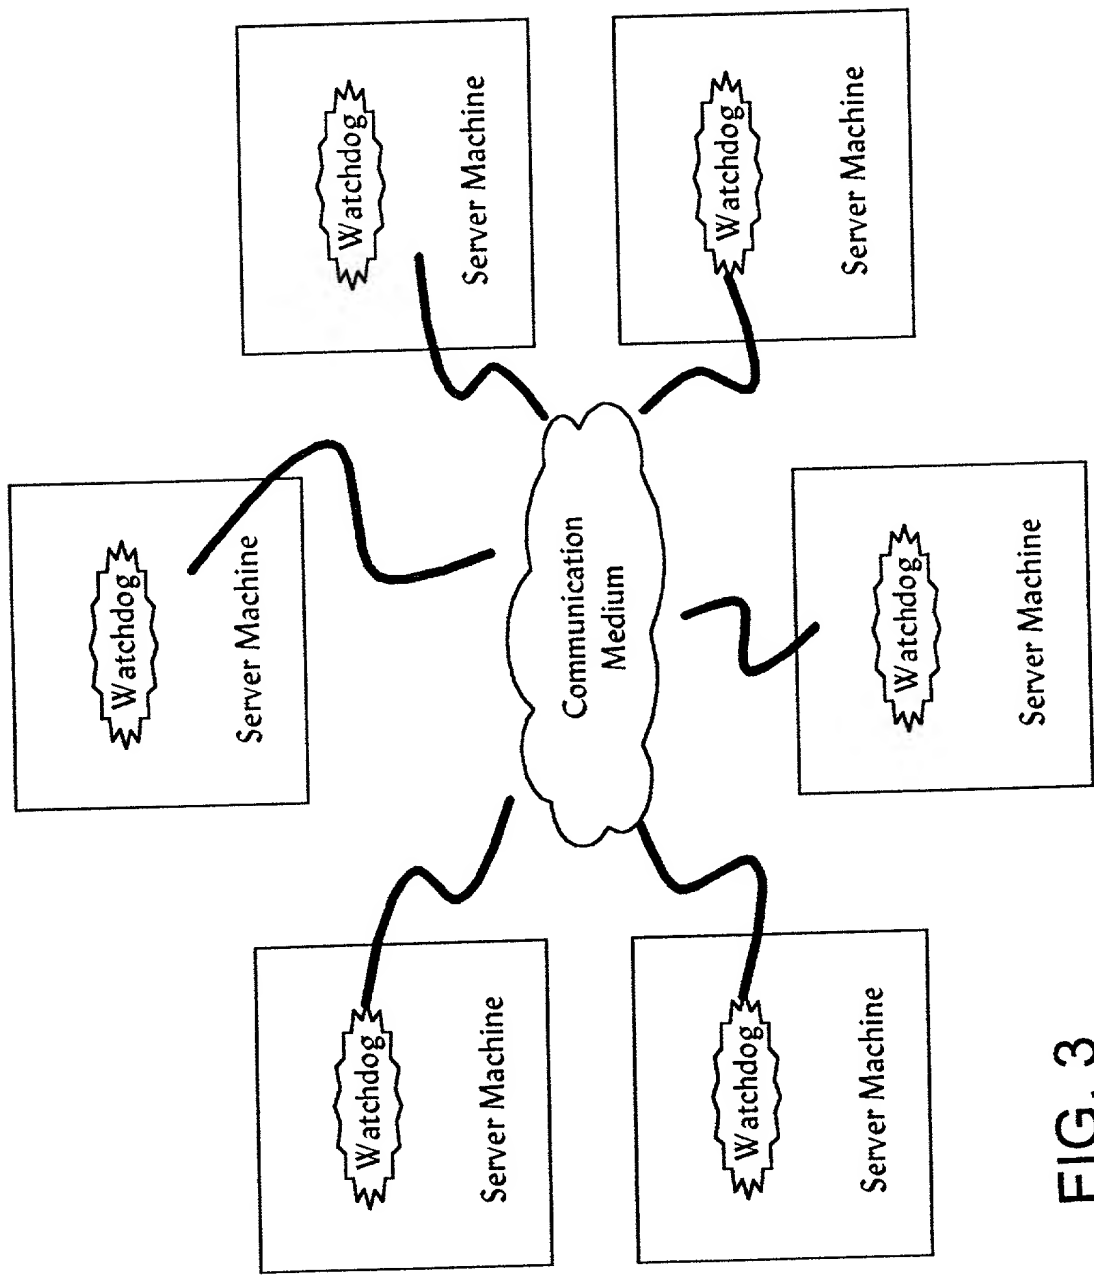


FIG. 3

Docket No.
DE919990052US1

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

IMPROVING AVAILABILITY AND SCALABILITY IN CLUSTERED APPLICATION SERVERS

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as United States Application No. or PCT International Application Number _____ and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)			Priority Not Claimed
99122914.7	Europe	18/11/99	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	
_____	_____	_____	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	
_____	_____	_____	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional application(s) listed below:

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. *(list name and registration number)*

Floyd A. Gonzalez, Reg. 26,732
 Lynn L. Augspurger, Reg. 24,227
 Lawrence D. Cutter, Reg. 28,501
 William A. Kinnaman, Jr., Reg. 27,650
 William B. Porter, Reg. 33,135
 Andrew J. Wojnicki, Jr., Reg.
 Christopher A. Hughes, Reg. 26,914
 Edward A. Pennington, Reg. 32,588
 John E. Hoel, Reg. 26,279
 Joseph C. Redmond, Jr., Reg. 18,753


Send Correspondence to: **Floyd A. Gonzalez**
IBM Corporation
2455 South Road, P386
Poughkeepsie, NY 12601

Direct Telephone Calls to: *(name and telephone number)*
Floyd A. Gonzalez, (845) 433-1163

Full name of sole or first inventor Frank Leymann	
Sole or first inventor's signature	Date
Residence Hasenaeckerweg 19, D-71134 Aidlingen, Federal Republic of Germany	
Citizenship German	
Post Office Address same as above	

Full name of second inventor, if any Dieter Roller	
Second inventor's signature	Date
Residence Hermann-Loens-Weg 5, D-71101 Schoenaich, Federal Republic of Germany	
Citizenship German	
Post Office Address same as above	

United States Patent & Trademark Office
Office of Initial Patent Examination -- Scanning Division



Application deficiencies were found during scanning:

☐ Page(s) _____ of _____ were not present:
for scanning. (Document title)

☐ Page(s) _____ of _____ were not present:
for scanning. (Document title)

Last page of the transmittal is misnumbered.

☐ Scanned copy is best available.

SCANNED, # 14